arXiv:1410.4066v1 [math.OC] 15 Oct 2014

Iteration Bounds for Finding the ϵ-Stationary Points for Structured Nonconvex Optimization

Bo Jiang * Shuzhong Zhang †

October 16, 2014

Abstract

In this paper we study proximal conditional-gradient (CG) and proximal gradient-projection type algorithms for a block-structured constrained nonconvex optimization model, which arises naturally from tensor data analysis. First, we introduce a new notion of ϵ-stationarity, which is suitable for the structured problem under consideration. We then propose two types of first-order algorithms for the model based on the proximal conditional-gradient (CG) method and the proximal gradient-projection method respectively. If the nonconvex objective function is in the form of mathematical expectation, we then discuss how to incorporate randomized sampling to avoid computing the expectations exactly. For the general block optimization model, the proximal subroutines are performed for each block according to either the block-coordinate-descent (BCD) or the maximum-block-improvement (MBI) updating rule. If the gradient of the nonconvex part of the objective f satisfies $\|\nabla f(x) - \nabla f(y)\|_q \leq M\|x - y\|_p^\delta$ where $\delta = p/q$ with $1/p + 1/q = 1$, then we prove that the new algorithms have an overall iteration complexity bound of $O(1/\epsilon^q)$ in finding an ϵ-stationary solution. If f is concave then the iteration complexity reduces to $O(1/\epsilon)$. Our numerical experiments for tensor approximation problems show promising performances of the new solution algorithms.

Keywords: constrained nonconvex optimization, block variables, iteration complexity bounds, conditional gradient algorithm

Mathematics Subject Classification: 90C26, 90C06, 90C15, 15A69.

*Research Center for Management Science and Data Analytics, School of Information Management and Engineering, Shanghai University of Finance and Economics, Shanghai 200433, China. Email: isyebojiang@gmail.com. Research of this author was supported in part by National Science Foundation of China (Grant 11401364).

†Department of Industrial and Systems Engineering, University of Minnesota, Minneapolis, MN 55455, USA. Email: zhangs@umn.edu. Research of this author was supported in part by National Science Foundation (Grant CMMI-1161242).

1 Introduction

The first-order algorithms and their iteration complexity analysis for nonconvex optimization problems have recently attracted considerable research attention; see e.g. [31, 7, 8, 9, 10, 18, 19]. In this paper, we aim to solve the following block-structured nonconvex and nonsmooth optimization problem:

$$\begin{array}{ll} \min & f(x_1, \cdots, x_d) + \sum_{i=1}^d h_i(x_i) \\ \text{s.t.} & x_i \in S_i \subseteq \mathbf{R}^{n_i}, \quad i = 1, \ldots, d, \end{array} \tag{1}$$

where f is differentiable but possibly nonconvex, and h_i is convex but possibly nonsmooth, for $i = 1, \ldots, d$. Such optimization models arise from a variety of applications. As an example, in statistics it is often desirable to find a *regulated* least square solution where the least square term is nonconvex, which falls into this category. Consider, for instance, the so-called sparse tensor PCA problem, where one wishes to find the best sparse rank-one approximation for a given tensor. Mathematically, this can be formulated by the following optimization model [1]:

$$\begin{array}{ll} \min & -\mathcal{A}(x_1, x_2, \cdots, x_d) + \rho \sum_{i=1}^d \|x_i\|_1 \\ \text{s.t.} & x_i \in S_i = \{x \,|\, \|x\|_2^2 \leq 1\}, \, i = 1, 2, ..., d, \end{array}$$

where $\mathcal{A}$ is a given d-dimensional tensor data. The L_1-norm in the objective is used to promote the sparsity of the vectors x_i, $i = 1, 2, ..., d$. The above problem is clearly an instance of (1).

Towards eventually solving (1), we first consider the special case when there is only one block of variables:

$$\begin{array}{ll} \min & \Phi(x) := f(x) + h(x) \\ \text{s.t.} & x \in S \subseteq \mathbf{R}^n. \end{array} \tag{2}$$

Again, here f is assumed to be differentiable but possibly nonconvex, and h is convex but possibly nonsmooth; S is assumed to be a compact and convex set, and Φ^* is the optimal value of (2). We denote

$$\mathrm{diam}_p(S) = \max_{x,y \in S} \|x - y\|_p, \tag{3}$$

where $\|x\|_p = (\sum_{i=1}^n |x_i|^p)^{1/p}$. In this paper, we shall propose two types of first-order algorithms to find a certain *ϵ-solution* for (2). A first question arises in this context: *Is there a reasonable definition of ϵ-solution (or rather, ϵ-stationary solution) to aim for?* For the smooth unconstrained version of (2), i.e. $\Phi(x) = f(x)$ and $S = \mathbf{R}^n$, the natural definition of ϵ-stationary point is:

$$\|\nabla \Phi(x)\|_2 = \|\nabla f(x)\|_2 \leq \epsilon.$$

Nesterov [31] and Cartis et al. [8] showed that the gradient decent type method with some properly chosen step size needs $O(1/\epsilon^2)$ iterations to find such a point. Moreover, Cartis et al. [7] constructed an example to show that the $O(1/\epsilon^2)$ complexity is actually tight for the steepest-descent algorithm.

However, the case for the constrained nonconvex optimization is more complicated. In fact, there are multiple quality measurements for approximative stationary points. Cartis et al. [9] proposed the following measure:

$$\chi_S(x) := \left| \min_{x+d\in S, \|d\|_2\leq 1} \nabla f(x)^\top d \right|.$$

Furthermore, they showed that it requires no more than $O(1/\epsilon^2)$ iterations for their adaptive cubic regularization algorithm to find a point x such that

$$\chi_S(x) \leq \epsilon. \tag{4}$$

Along a related but different line, Ghadimi et al. [18] used the squared norm of the residual for some generalized projection to evaluate the quality of solution. Specifically, the residual at point x is defined as

$$P_S(x,\gamma) := \frac{1}{\gamma}(x - x^+),$$

where

$$x^+ = \arg\min_{y\in S} \nabla f(x)^\top y + \frac{1}{\gamma} V(y,x) + h(y),$$

and V is some *prox-function*. (We refer the interested reader to [18] for the details). The authors proposed a projected gradient algorithm and proved that it will take no more than $O(1/\epsilon^2)$ iterations to achieve

$$\|P_S(x,\gamma)\|_2^2 \leq \epsilon. \tag{5}$$

In this paper we consider the following new notion of stationarity for the nonconvex and nondifferentiable optimization model (2):

Definition 1.1 *We call x to be a stationary point of (2) if the following condition holds:*

$$\nabla f(x)^\top (y - x) + h(y) - h(x) \geq 0 \quad \forall y \in S. \tag{6}$$

In fact, if x is a local minimizer of (2), then it must satisfy (6). To see this, we shall use a contradiction argument. Suppose that there exists some $y \in S$ such that $\nabla f(x)^\top (y - x) + h(y) - h(x) < 0$. Denote $d = y - x$. Then the directional derivative along direction d at point x satisfies

$$\begin{aligned}
(f+h)'(x;d) &= \lim_{\alpha\downarrow 0} \frac{f(x+\alpha d) + h(x+\alpha d) - f(x) - h(x)}{\alpha} \\
&\leq \lim_{\alpha\downarrow 0} \frac{f(x+\alpha d) - f(x)}{\alpha} + \lim_{\alpha\downarrow 0} \frac{(1-\alpha)h(x) + \alpha h(x+d) - h(x)}{\alpha} \\
&= \nabla f(x)^\top (y - x) + h(y) - h(x) < 0,
\end{aligned}$$

where the inequality is due to the convexity of h. Consequently, x cannot be a local optimal solution of problem (2). Thus based on condition (6), we consider the following definition of approximative stationary solution.

Definition 1.2 *We call x to be an ϵ-stationary point of (2) if*

$$\psi_S(x) := \nabla f(x)^\top (y - x) + h(y) - h(x) \geq -\epsilon \quad \forall y \in S. \tag{7}$$

A similar condition for L_2-L_p optimization problem was considered in [17]; the relationship between conditions (7), (4) and (5) will be discussed in Section 2. To proceed, let us make the following technical assumption on the smooth part of the objective $f(x)$ throughout this paper.

Assumption 1.1 *There exists some $p > 1$ and $\lambda > 0$ such that*

$$f(y) \leq f(x) + \nabla f(x)^\top (y - x) + \frac{\lambda}{2}\|y - x\|_p^p, \ \ \forall x, y \in S. \tag{8}$$

Some comments about Assumption 1.1 are in order here. First, notice that if $f(x)$ is concave, then (8) holds true for any $p > 0$ and $\lambda > 0$. Second, if the gradient of f satisfies

$$\|\nabla f(x) - \nabla f(y)\|_q^q \leq M\|x - y\|_p^p \quad \forall\, x, y \in S \tag{9}$$

for some $p, q > 1$ and $\frac{1}{p} + \frac{1}{q} = 1$, then the function itself also satisfies (8). To see this, we let $z = y - x$ and $g(\alpha) = f(x + \alpha z)$. It follows that

$$\begin{aligned}
f(y) - f(x) &= \int_0^1 g'(\alpha) z \, d\alpha = \int_0^1 \nabla f(x + \alpha z)^\top z \, d\alpha \\
&\leq \int_0^1 \nabla f(x)^\top z \, d\alpha + \left| \int_0^1 (\nabla f(x + \alpha z) - \nabla f(x))^\top z \, d\alpha \right| \\
&\leq \int_0^1 \nabla f(x)^\top z \, d\alpha + \int_0^1 \|\nabla f(x + \alpha z) - \nabla f(x)\|_q \|z\|_p \, d\alpha \\
&\leq \nabla f(x)^\top z + M^{1/q} \|z\|_p^{1 + \frac{p}{q}} \int_0^1 \alpha^{\frac{p}{q}} \, d\alpha \\
&= \nabla f(x)^\top z + \frac{M^{1/q}}{p} \|z\|_p^p,
\end{aligned}$$

where the last equality follows from $\frac{1}{p} + \frac{1}{q} = 1$. Thus, the function with Lipschitz continuous gradient automatically satisfies inequality (8) for $p = q = 2$. In fact, condition (9) reflects the degree of the Hölderian continuity of ∇f, which was also considered in [13] to construct an inexact first order oracle. Finally, we remark that the p-th powered p-norm function:

$$f(x) = \sum_{i=1}^n x_i^p, \ \forall\, 1 < p \leq 2, \tag{10}$$

on $\mathbf{R}_+^n$ also satisfies (9). We observe that the function is separable with respect to all x_i, and so it suffices to show that there exists some λ such that:

$$v^p \leq u^p + (u^p)'(v - u) + \frac{\lambda}{2}|v - u|^p = u^p + p\, u^{p-1}(v - u) + \frac{\lambda}{2}|v - u|^p, \tag{11}$$

when $1 < p \leq 2$. If $u = 0$, then the inequality trivially holds for any $\lambda \geq 2$; otherwise we can divide both sides by $|u|^p$ and aim to prove an equivalent formulation:

$$k^p \leq 1 + p(k-1) + \frac{\lambda}{2}|k-1|^p, \tag{12}$$

where $k = v/u$. To this end, define

$$g(k) := \begin{cases} 0, & \text{if } k = 1 \\ \frac{k^p - 1 - p(k-1)}{|k-1|^p}, & \text{otherwise.} \end{cases}$$

Observe that $\lim_{k \to +\infty} g(k) = 1$, and from L'Hospital's rule

$$\lim_{k \to 1} g(k) = \begin{cases} 0, & \text{if } 1 < p < 2 \\ 1, & \text{if } p = 2, \end{cases}$$

and so $g(k)$ is upper bounded on $\mathbf{R}$ and there exits some $\hat{\lambda}$ such that (12) holds. Finally by letting $\lambda = \max\{2, \hat{\lambda}\}$, the inequality (11) follows.

In this paper we shall propose two algorithms for solving problem (2), both achieving an $O(1/\epsilon^q)$ iteration complexity, where $\frac{1}{p} + \frac{1}{q} = 1$ and p is the parameter in (8). As a result, a larger value of p leads to a smaller value of q, hence a better iteration bound for the algorithm. In other words, this result shows that the "smoothness" of the function will be reflected in the speed of the convergence. In particular, when $p = 2$, we get $O(1/\epsilon^q) = O(1/\epsilon^2)$, which is consistent with the result of Cartis et al. [9]. Another extreme case is when $f(x)$ is concave, and in this case the complexity can be reduced to $O(1/\epsilon)$; we shall elaborate more on this point later.

The algorithms to be proposed in this paper use only the first-order information of f. When $p = 2$ and h does not appear, **Algorithm 2** in this paper is simply the gradient projection algorithm. When $p = 2$ and f is *convex*, **Algorithm 2** coincides with the so-called ISTA (iterative shrinkage-thresholding algorithm ([3]). Similarly, if h does not appear, then **Algorithm 1** coincides with the conditional gradient (CG) method, where the subproblem to be solved in each iteration involves a linear objective function. Note that the conditional gradient method was originally proposed by Frank and Wolfe [15], and recently has regained some research attention primarily due to the fact that the linear subproblem is easier to solve in the context of large scale optimization. To the best of our knowledge, the iteration complexity analysis for the CG method had only been established for convex optimization [16, 26]. In other words, the current paper presents for the first time an iteration complexity bound for the CG method in the context of nonconvex optimization. A recent work related to the current paper is [28], which proposes a smoothing SQP method to solve (2). The key differences are: (1) In [28], $f(x)$ is assumed to be in the form of $\|(Ax-b)_+\|_q^q$ where $0 < q < 1$, and the gradient of $h(x)$ is assumed to be Lipschitz continuous and h may also be non-convex; (2) In [28] the constraint set S is assumed to be polyhedral; (3) In [28] a convex quadratic program is solved at each step as a subroutine. As we can see, the basic assumptions on the problem setting as well as the subroutines applied are all very different. The results are

also fundamentally different. In [28], a different notion of ϵ-KKT condition is introduced, and a smoothing SQP method is shown to reach an ϵ-KKT point in no more than $O(\epsilon^{q-4})$ iterations. In the current paper, the ϵ-stationarity condition is based on a variational inequality. Though $h(x)$ is assumed to be convex in our context, the CG subroutine may be much easier to solve. In this paper, we also extend our studies to stochastic nonconvex optimization, for which a combination of randomized sampling method and the first-order approximation approach is proposed. We also consider the case where the nonsmooth part is concave and show that the CG method achieves similar complexity bound by properly incorporating some randomized smoothing scheme. Besides, we show that our approach can be modified to handle multi-block nonconvex optimization which covers a great variety of applications.

This paper is organized as follows. In Section 2 we introduce a new notion of ϵ-stationary solution and discuss its relationships to those proposed by Cartis et al. and Ghadimi et al.. In Section 3, we present two algorithms for problem (2) and analyze their iteration complexity bounds. We then develop a stochastic algorithm and a randomized smoothing algorithm in Section 4. In Section 5, the solution methods are extended to solve a nonconvex multi-block optimization model. Finally, we present our numerical experiments in Section 6.

2 The ϵ-Stationarity Condition

2.1 Relationship with Existing Results

In the last section, we introduced two quality measures for an approximate solution, denoted by $\chi_S(x)$ and $\|P_S(x,\gamma)\|_2^2$ respectively. We also proposed our new quality measure $\psi_S(x)$ (formula (7)). Obviously, these three measures are different, but they are related. To be precise, their relationship is summarized in the following proposition.

Proposition 2.1 *(i) If $h(x) \equiv 0$ and $\psi_S(x) \geq -\epsilon$, then $\chi_S(x) \leq \epsilon$;*
(ii) Suppose the prox-function $V(y,x) = \|y-x\|_2^2/2$, then $\psi_S(x) \geq -\epsilon$ implies $\|P_S(x,\gamma)\|_2^2 \leq \frac{\epsilon}{\gamma}$. Conversely if we further assume the gradient function $\nabla f(x)$ is continuous, then $\|P_S(x,\gamma)\|_2^2 \leq \epsilon$ implies

$$\psi_S(x) \geq -(\gamma\,\tau + \gamma\,\varsigma + \mathrm{diam}_2(S))\sqrt{\epsilon},$$

where $\tau = \max_{x\in S} \|\nabla f(x)\|_2$, $\varsigma = \max_{x\in S} \min_{z\in\partial h(x)} \|z\|_2$ and $\mathrm{diam}_2(S)$ is defined in (3).

Proof. According to the definition,

$$\begin{aligned} & \chi_S(x) \leq \epsilon \\ \Longleftrightarrow \quad & \min_{x+d\in S, \|d\|_2\leq 1} \nabla f(x)^\top d \geq -\epsilon \\ \Longleftrightarrow \quad & \nabla f(x)^\top (y-x) \geq -\epsilon,\ \forall\ \|y-x\|_2 \leq 1,\ y \in S. \end{aligned}$$

Therefore, (i) is readily implied by (7) when $h(x) \equiv 0$.
Now, let us prove (ii). Since $V(y,x) = \|y-x\|_2^2/2$, one has

$$\left(\nabla f(x) + \frac{1}{\gamma}(x^+ - x) + z\right)^\top (y - x^+) \geq 0 \quad \forall\, y \in S, \tag{13}$$

where $z \in \partial h(x^+)$. We can particularly choose $y = x$ and get

$$\nabla f(x)^\top (x - x^+) + h(x) - h(x^+) \geq (\nabla f(x) + z)^\top (x - x^+) \geq \frac{1}{\gamma}\|x^+ - x\|_2^2.$$

So if $\psi_S(x) \geq -\epsilon$ (i.e. (7) holds) then we have $\|P_S(x,\gamma)\|_2^2 \leq \frac{\epsilon}{\gamma}$. To show the other direction, note that ς is finite (S is compact) and $h(x)$ is convex. Thus, for $x, x^+ \in S$ we can choose $w \in \partial h(x)$ such that

$$\varsigma \, \|x - x^+\|_2 \geq w^\top (x - x^+) \geq h(x) - h(x^+).$$

This inequality together with (13) implies that

$$\begin{aligned}
& \nabla f(x)^\top (y - x) + h(y) - h(x) + (\|\nabla f(x)\|_2 + \varsigma)\|x - x^+\|_2 \\
\geq\ & \nabla f(x)^\top (y - x) + h(y) - h(x) + \nabla f(x)^\top (x - x^+) + h(x) - h(x^+) \\
=\ & \nabla f(x)^\top (y - x^+) + h(y) - h(x^+) \\
\geq\ & (\nabla f(x) + z)^\top (y - x^+) \\
\geq\ & -\frac{1}{\gamma}(x^+ - x)^\top (y - x^+) \\
\geq\ & -\frac{1}{\gamma}\|y - x^+\|_2 \|x^+ - x\|_2 \quad \forall\, y \in S,
\end{aligned}$$

where $z \in \partial h(x^+)$ and the second inequality follows from the convexity of $h(x)$. Furthermore, since S is compact and $\nabla f(x)$ is continuous, by rearranging the terms in the above inequality, the final conclusion follows. □

Under the conditions of Proposition 2.1, the relationship among these three measures for an approximate local optimal solution is depicted in Figure 1. According to this result, we see that condition (7) is in some sense more general than (4) and (5). In the remainder of this paper, we shall refer the ϵ-stationary condition to (7).

2.2 Sufficient Conditions for ϵ-Stationarity

For a given point z, we define the following two functions, which will play a crucial role in our solution methods to be proposed later:

$$L(x;z) := f(z) + \nabla f(z)^\top (x - z) + h(x), \tag{14}$$

$$U(x;z) := f(z) + \nabla f(z)^\top (x - z) + \frac{\lambda}{2}\|x - z\|_p^p + h(x). \tag{15}$$

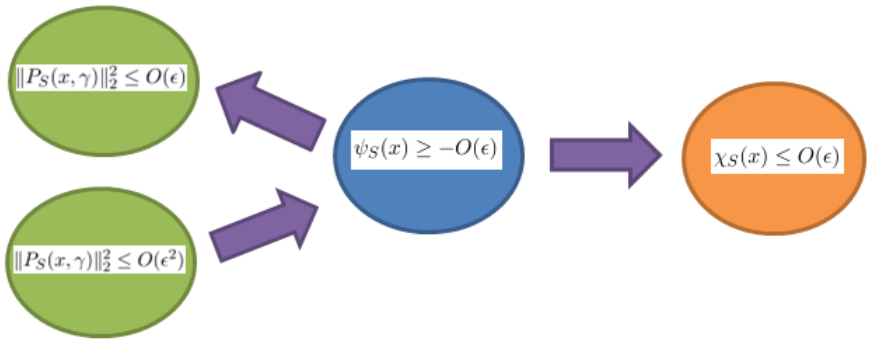

Figure 1: Relationship of the three different ϵ-stationarity measures.

In fact, (14) is obtained by linearizing the smooth part of Φ and (15) is an upper bound of Φ if Assumption 1.1 holds. These two functions lead to the following two convex optimization subroutines:

$$\begin{cases} \min\limits_{x} & L(x;z) \\ \text{s.t.} & x \in S, \end{cases} \tag{16}$$

and

$$\begin{cases} \min\limits_{x} & U(x;z) \\ \text{s.t.} & x \in S. \end{cases} \tag{17}$$

Denote z_L and z_U to be the minimizer of (16) and (17) respectively. The partial linearized improvement at point z is defined by

$$\triangle L_z := L(z;z) - L(z_L;z) = -\nabla f(z)^\top (z_L - z) + h(z) - h(z_L).$$

Similarly we define the partial p-powered improvement by

$$\triangle U_z := U(z;z) - U(z_U;z) = -\nabla f(z)^\top (z_U - z) - \frac{\lambda}{2}\|z_U - z\|_p^p + h(z) - h(z_U).$$

The following lemma, which is inspired by [17], states that if the progress gained by solving (16) or (17) is small, then we are already near a stationary point.

Lemma 2.2 *Given $\epsilon \ge 0$, for any $z \in S$,*
(i) if $\triangle L_z \le \epsilon$, then z is an ϵ-stationary point of (2);
(ii) if $\triangle U_z \le \frac{1}{2}\left(\frac{\epsilon}{\mathrm{diam}_p(S)\lambda^{1/p}}\right)^q$ with $\frac{1}{p}+\frac{1}{q}=1$, and $\epsilon \le \mathrm{diam}_p^p(S)\lambda$, then z is an ϵ-stationary point of (2).

Proof. Let us first consider (i). Since z_L is optimal to (16), we have

$$L(y;z) - L(z_L;z) = \nabla f(z)^\top (y - z_L) + h(y) - h(z_L) \ge 0, \ \forall\, y \in S.$$

It follows that

$$\begin{aligned} & \nabla f(z)^\top (y-z) + h(y) - h(z) \\ = \; & \nabla f(z)^\top (y-z_L) + h(y) - h(z_L) + \nabla f(z)^\top (z_L - z) + h(z_L) - h(z) \\ \geq \; & \nabla f(z)^\top (z_L - z) + h(z_L) - h(z) \quad \forall\, y \in S. \end{aligned}$$

Then by definition, $\triangle L_z \leq \epsilon$ implies that

$$\nabla f(z)^\top (y-z) + h(y) - h(z) \geq -\triangle L_z \geq -\epsilon.$$

To prove statement (ii), we consider the point $y = z + s(z_L - z)$ with $0 \leq s \leq 1$. Clearly, by definition of partial p-powered improvement and convexity of h, it follows that

$$\begin{aligned} \frac{1}{2}\left(\frac{\epsilon}{\mathrm{diam}_p(S)\lambda^{1/p}}\right)^q & \geq \triangle U_z \geq U(z;z) - U(y;z) \\ & = \nabla f(z)^\top (z-y) - \frac{\lambda}{2}\|y-z\|_p^p + h(z) - h(y) \\ & \geq -s\nabla f(z)^\top (z_L - z) - \frac{\lambda}{2} s^p \|z_L - z\|_p^p + s\left(h(z) - h(z_L)\right). \end{aligned}$$

Letting $s = \left(\frac{\epsilon}{\mathrm{diam}_p^p(S)\lambda}\right)^{\frac{1}{p-1}}$ and rearranging the terms in the above inequality yield

$$\begin{aligned} \triangle L_z & = -\nabla f(z)^\top (z_L - z) + h(z) - h(z_L) \\ & \leq \frac{\lambda}{2} s^{p-1} \|z_L - z\|_p^p + \frac{1}{2s}\left(\frac{\epsilon}{\mathrm{diam}_p(S)\lambda^{1/p}}\right)^q \\ & = \frac{\lambda}{2}\frac{\epsilon}{\mathrm{diam}_p^p(S)\lambda}\|z_L - z\|_p^p + \frac{1}{2}(\mathrm{diam}_p^p(S)\lambda)^{\frac{1}{p-1}-\frac{q}{p}}\epsilon^{q-\frac{1}{p-1}}. \end{aligned}$$

Since $\frac{1}{p} + \frac{1}{q} = 1$, we have $\frac{1}{p-1} - \frac{q}{p} = 0$ and $q - \frac{1}{p-1} = 1$. These facts together with the definition of $\mathrm{diam}_p(S)$, as well as the inequality above, imply that

$$\triangle L_z \leq \frac{\epsilon}{2} + \frac{\epsilon}{2} = \epsilon,$$

which, combined with statement (i), proves the desired result. □

3 Algorithms and Their Iteration Complexities for Finding the ϵ-Stationary Point

We are now in a position to present our first algorithm for (2). In particular, at each iteration we find the search direction through optimizing a partially linearized function and then determine the step size by minimizing a simple one-dimensional function.

Algorithm 1

Let $x^0 \in S$ be given and set $y^0 = x^0$.

for $k = 1, 2, \cdots, N$, **do**

$y^k = \arg\min_{y \in S} L(y; x^k)$, and let $d^k = y^k - x^k$;

$\alpha_k = \arg\min_{\alpha \in [0,1]} \alpha \nabla f(x^k)^\top d^k + \alpha^p \frac{\lambda}{2} \|d^k\|_p^p + (1-\alpha) h(x^k) + \alpha\, h(y^k)$.

Set $x^{k+1} = (1-\alpha_k) x^k + \alpha_k y^k$.

end for

Note that in the absence of the nonsmooth part h, this algorithm is simply CG (Conditional Gradient). The computational complexity of **Algorithm 1** to reach an ϵ-stationary solution is as follows.

Theorem 3.1 *For any* $0 < \epsilon < \mathrm{diam}_p^p(S)\lambda$, **Algorithm 1** *finds an* ϵ*-stationary point of* (2) *within* $\left\lceil \frac{2(\Phi(x^1) - \Phi^*)(\mathrm{diam}_p^p(S)\lambda)^{q-1}}{\epsilon^q} \right\rceil$ *steps, where* $\frac{1}{p} + \frac{1}{q} = 1$.

Proof. According to Assumption 1.1, it holds that

$$\begin{aligned} & -\nabla f(x^k)^\top (x^{k+1} - x^k) - \frac{\lambda}{2} \|x^{k+1} - x^k\|_p^p + h(x^k) - h(x^{k+1}) \\ \leq\ & f(x^k) - f(x^{k+1}) + h(x^k) - h(x^{k+1}) \\ =\ & \Phi(x^k) - \Phi(x^{k+1}). \end{aligned} \tag{18}$$

Note that $\frac{\epsilon}{\mathrm{diam}_p^p(S)\lambda} \leq 1$ and $x^{k+1} - x^k = \alpha_k (y^k - x^k)$. For simplicity, denote $\triangle L_k := \triangle L_{x^k}$. By the optimality of α_k, we have

$$\begin{aligned} & \left(\frac{\epsilon}{\mathrm{diam}_p^p(S)\lambda}\right)^{\frac{1}{p-1}} \triangle L_k - \frac{1}{2\lambda^{1/(p-1)}} \left(\frac{\epsilon}{\mathrm{diam}_p(S)}\right)^{\frac{p}{p-1}} \\ \leq\ & -\left(\frac{\epsilon}{\mathrm{diam}_p^p(S)\lambda}\right)^{\frac{1}{p-1}} \Big(\nabla f(x^k)^\top (y^k - x^k) + h(y^k) - h(x^k)\Big) - \frac{\lambda}{2} \frac{\|y^k - x^k\|_p^p}{\mathrm{diam}_p^p(S)} \left(\frac{\epsilon}{\lambda \mathrm{diam}_p(S)}\right)^{\frac{p}{p-1}} \\ =\ & -\left(\frac{\epsilon}{\mathrm{diam}_p^p(S)\lambda}\right)^{\frac{1}{p-1}} \Big(\nabla f(x^k)^\top (y^k - x^k) + h(y^k) - h(x^k)\Big) - \frac{\lambda}{2} \left(\frac{\epsilon}{\mathrm{diam}_p^p(S)\lambda}\right)^{\frac{p}{p-1}} \|y^k - x^k\|_p^p \\ \leq\ & -\alpha_k \Big(\nabla f(x^k)^\top (y^k - x^k) + h(y^k) - h(x^k)\Big) - \frac{\lambda \alpha_k^p}{2} \|y^k - x^k\|_p^p \\ =\ & -\nabla f(x^k)^\top (\alpha_k (y^k - x^k)) + h(x^k) - (1-\alpha_k) h(x^k) - \alpha_k h(y^k) - \frac{\lambda}{2} \|\alpha_k (y^k - x^k)\|_p^p \\ \leq\ & -\nabla f(x^k)^\top (x^{k+1} - x^k) + h(x^k) - h(x^{k+1}) - \frac{\lambda}{2} \|x^{k+1} - x^k\|_p^p, \end{aligned}$$

where the last inequality is due to the convexity of function $h(\cdot)$. Combining this formula with (18) leads to

$$\left(\frac{\epsilon}{\mathrm{diam}_p^p(S)\lambda}\right)^{\frac{1}{p-1}} \triangle L_k \leq \Phi(x^k) - \Phi(x^{k+1}) + \frac{1}{2\lambda^{1/(p-1)}} \left(\frac{\epsilon}{\mathrm{diam}_p(S)}\right)^{\frac{p}{p-1}}.$$

Dividing both sides by $\left(\frac{\epsilon}{\text{diam}_p^p(S)\lambda}\right)^{\frac{1}{p-1}}$, one has

$$\triangle L_k \leq \left(\frac{\epsilon}{\text{diam}_p^p(S)\lambda}\right)^{-\frac{1}{p-1}} \left(\Phi(x^k) - \Phi(x^{k+1})\right) + \frac{\epsilon}{2}.$$

Summing up the above inequalities for $k = 1, \ldots, N$ yields

$$\begin{aligned} N \min_{k\in\{1,\ldots,N\}} \triangle L_k &\leq \sum_{k=1}^{N} \triangle L_k \leq \left(\frac{\epsilon}{\text{diam}_p^p(S)\lambda}\right)^{-\frac{1}{p-1}} \left(\Phi(x^1) - \Phi(x^{N+1})\right) + \frac{\epsilon}{2}N \\ &\leq \left(\frac{\epsilon}{\text{diam}_p^p(S)\lambda}\right)^{-\frac{1}{p-1}} \left(\Phi(x^1) - \Phi^*\right) + \frac{\epsilon}{2}N. \end{aligned}$$

Observe that $\frac{1}{p} + \frac{1}{q} = 1$ leads to $q - 1 - \frac{1}{p-1} = 0$. Dividing the above inequalities by $N = \left\lceil \frac{2(\Phi(x^1)-\Phi^*)(\text{diam}_p^p(S)\lambda)^{q-1}}{\epsilon^q} \right\rceil$, we arrive at the conclusion that there must exist some $\tilde{k} \leq N$ such that $\triangle L_{\tilde{k}} \leq \epsilon$, which combined with Lemma 2.2 further implies that $x^{\tilde{k}}$ is an ϵ-stationary point of (2). □

From Theorem 3.1, we can see that the larger value of p implies fewer iteration numbers required by the algorithm. Since a concave function satisfies Assumption 1.1 for any $p \geq 1$, in this case a better complexity bound is guaranteed.

Corollary 3.2 *If f is a concave function, then performing* **Algorithm 1** *with $\alpha_k = 1$ for all k will reach an ϵ-stationary point of (2) within $\left\lceil \frac{\Phi(x^1)-\Phi^*}{\epsilon} \right\rceil$ steps.*

To proceed, we present our second algorithm below, which is based on minimizing an upper bound of the original objective function at each iteration.

Algorithm 2

Let $x^1 \in S$ be given

for $k = 1, 2, \cdots, N$, **do**

$x^{k+1} = \arg\min_{y\in S} U(y; x^k)$.

end for

We remark that this algorithm still solves a separable convex subproblem since the function $U(y; x^k)$ itself is convex when $p \geq 1$ and is separable with respect to y_i for all i. For simplicity, we denote $\triangle U_k := \triangle U_{x^k,p}$. Below is the complexity result for **Algorithm 2**.

Theorem 3.3 *For any any $0 < \epsilon < \text{diam}_p^p(S)\lambda$,* **Algorithm 2** *finds an ϵ-stationary point of (2) within $\left\lceil \frac{2(\Phi(x^1)-\Phi^*)(\text{diam}_p^p(S)\lambda)^{q-1}}{\epsilon^q} \right\rceil$ steps, where $\frac{1}{p} + \frac{1}{q} = 1$.*

Proof. According to Assumption 1.1, it holds that

$$f(x^{k+1}) \leq f(x^k) + \nabla f(x^k)^\top (x^{k+1} - x^k) + \frac{\lambda}{2}\|x^{k+1} - x^k\|_p^p.$$

As a result,

$$\begin{aligned} \triangle U_k &= -\nabla f(x^k)^\top (x^{k+1} - x^k) - \frac{\lambda}{2}\|x^{k+1} - x^k\|_p^p + h(x^k) - h(x^{k+1}) \\ &\leq -f(x^{k+1}) + f(x^k) + h(x^k) - h(x^{k+1}) \\ &= \Phi(x^k) - \Phi(x^{k+1}). \end{aligned}$$

Summing up the above inequalities for $k = 1, \ldots, N$ yields

$$N \min_{k\in\{1,\ldots,N\}} \triangle U_k \leq \sum_{k=1}^{N} \triangle U_k \leq \Phi(x^1) - \Phi(x^{N+1}) \leq \Phi(x^1) - \Phi^*.$$

Thus, setting $N = \left\lceil \frac{2(\Phi(x^1)-\Phi^*)(\mathrm{diam}_p^p(S)\lambda)^{q-1}}{\epsilon^q} \right\rceil$, there must exist some $\tilde{k} \leq N$ such that

$$\triangle U_{\tilde{k}} \leq \frac{1}{2(\mathrm{diam}_p^p(S)\lambda)^{q-1}}\epsilon^q = \frac{1}{2}\left(\frac{\epsilon}{\mathrm{diam}_p(S)\lambda^{1/p}}\right)^q.$$

This inequality combined with statement (ii) in Lemma 2.2 implies that $x^{\tilde{k}}$ is an ϵ-stationary point for (2). □

4 Iteration Complexity Bounds for Stochastic and Smoothing Approximation Methods

4.1 Complexity for Stochastic Approximation

In this subsection, we study the case where the exact gradient of $f(x)$ in problem (2) is not available. Instead, we assume that a noise estimation for the gradient of f can be obtained by resorting to the so-called stochastic first-order oracle $\mathcal{SFO}$. In particular, for point $x^k \in S$ at the k-th iteration, $\mathcal{SFO}$ would return a stochastic gradient $G(x^k, \xi^k)$, where ξ^k is a random variable satisfying

$$\mathsf{E}[G(x^k, \xi^k)] = \nabla f(x^k), \tag{19}$$

$$\mathsf{E}[\|G(x^k, \xi^k) - \nabla f(x^k)\|_q^q] \leq \sigma^q, \tag{20}$$

for some constant $\sigma > 0$.

The method of randomized sampling for stochastic programming can be traced back to the seminal paper of Robbins and Monro [35] (1951). Computational complexity for convex optimization

was first studied in [30]; similar results were subsequently established for convex stochastic optimization [29, 25]. Recently, complexity analysis for stochastic approximation has been successfully extended to several nonconvex models as well; cf. [19, 18, 38].

Inspired by [18], we propose below a mini-batch stochastic algorithm for the stochastic version of (2):

Algorithm 3

Let $x^1 \in S$ be given

for $k = 1, 2, \cdots, N$, **do**

Call the $\mathcal{SFO}$ m_k times to obtain $G(x^k, \xi^{k,i})$, $i = 1, \ldots, m_k$;

set $G_k = \frac{1}{m_k} \sum_{i=1}^{m_k} G(x^k, \xi^{k,i})$, and compute

$x^{k+1} = \arg\min_{x \in S} \tilde{U}(x; G_k)$, where $\tilde{U}(x; G_k) := G_k^\top (x - x^k) + \frac{\lambda}{2}\|x - x^k\|_p^p + h(x)$.

end for

Before discussing the computational complexity of **Algorithm 3**, we shall note the following two technical lemmas.

Lemma 4.1 *For any $a, b, c \in \mathbf{R}$ and $p \geq 2$ we have*

$$(\mathrm{sign}(a-c)|a-c|^{p-1} - \mathrm{sign}(b-c)|b-c|^{p-1})(a-b) \geq (1/2)^{p-2}|a-b|^p. \tag{21}$$

Proof. First of all, we observe that x^{p-1} is a convex function for $x \in \mathbf{R}_+$, and so for any $x, y \in \mathbf{R}_+$ we have

$$\left(\frac{x+y}{2}\right)^{p-1} \leq \frac{1}{2}\left(x^{p-1} + y^{p-1}\right)$$

implying that $x^{p-1} + y^{p-1} \geq \frac{1}{2^{p-2}}(x+y)^{p-1}$. Also, we have $(x+y)^{p-1} \geq x^{p-1} + y^{p-1}$, because $(x^{p-1} + y^{p-1})^{1/(p-1)}$ is the L_{p-1}-norm of (x, y) which can never exceed its corresponding L_1-norm.

To prove the lemma, due to symmetry we need only to consider three separate cases: (i) $c < b < a$; (ii) $c < a < b$; (iii) $a < c < b$.

In case (i), the LHS of (21) equals

$$\left[(a-c)^{p-1} - (b-c)^{p-1}\right](a-b) \geq (a-c)^{p-1}(a-b) \geq (a-b)^p \geq \frac{1}{2^{p-2}}(a-b)^p.$$

In case (ii), the LHS of (21) equals

$$\left[(b-c)^{p-1} - (a-c)^{p-1}\right](b-a) \geq (b-a)^{p-1}(b-a) \geq \frac{1}{2^{p-2}}(a-b)^p.$$

Finally, in case (iii), the LHS of (21) equals

$$\left[(c-a)^{p-1} + (b-c)^{p-1}\right](b-a) \geq \frac{1}{2^{p-2}}(b-a)^{p-1}(b-a) = \frac{1}{2^{p-2}}(b-a)^p.$$

Summarizing all the cases, the claimed inequality (21) follows. □

Lemma 4.2 *For integer $p \geq 2$, let*

$$x_i^* = \min_{x \in S} g_i^\top (x - z) + \frac{\lambda}{2}\|x - z\|_p^p + h(x), \ \ \textit{for } i = 1, 2.$$

Then it holds that

$$(1/2)^p \|x_1^* - x_2^*\|_p^p \leq 1/(\lambda p)^q \|g_1 - g_2\|_q^q, \ \ \textit{with } 1/p + 1/q = 1.$$

Proof. By the optimality of x_1^* and x_2^*, for any $x \in S$, there exist $w_1 \in \partial h(x_1^*)$ and $w_2 \in \partial h(x_2^*)$ such that

$$\left(g_1 + \frac{\lambda}{2} p \,\mathrm{sign}(x_1^* - z)\|x_1^* - z\|_{p-1}^{p-1} + w_1\right)^\top (x - x_1^*) \geq 0, \tag{22}$$

and

$$\left(g_2 + \frac{\lambda}{2} p \,\mathrm{sign}(x_2^* - z)\|x_2^* - z\|_{p-1}^{p-1} + w_2\right)^\top (x - x_2^*) \geq 0. \tag{23}$$

Letting $x = x_2^*$ in (22) and $x = x_1^*$ in (23), by Lemma 4.1 and summing up (22) and (23), we have

$$\begin{aligned}
& (g_1 - g_2)^\top (x_2^* - x_1^*) \\
\geq \ & \frac{\lambda}{2} p \left(\mathrm{sign}(x_1^* - z)\|x_1^* - z\|_{p-1}^{p-1} - \mathrm{sign}(x_2^* - z)\|x_2^* - z\|_{p-1}^{p-1}\right)^\top (x_1^* - x_2^*) + (w_1 - w_2)^\top (x_1^* - x_2^*) \\
\geq \ & \frac{\lambda\, p}{2^{p-1}} \|x_1^* - x_2^*\|_p^p + (w_1 - w_2)^\top (x_1^* - x_2^*) \\
\geq \ & \frac{\lambda\, p}{2^{p-1}} \|x_1^* - x_2^*\|_p^p,
\end{aligned}$$

where the last inequality is due to the convexity of h. On the other hand, by the Hölder inequality one has

$$(g_1 - g_2)^\top (x_2^* - x_1^*) \leq \|g_1 - g_2\|_q \|x_1^* - x_2^*\|_p \ \text{ with } 1/p + 1/q = 1.$$

The desired result follows by combining these two inequalities and then taking the q-th power on both sides. □

Recall that $\tilde{U}(x; G_k) = G_k^\top (x - x^k) + \frac{\lambda}{2}\|x - x^k\|_p^p + h(x)$. We are ready to present the main result of this subsection.

Theorem 4.3 *Suppose $\{x^k\}$ is the sequence of iterates generated by* **Algorithm 3**. *Denote*

$$\triangle \tilde{U}_k = \tilde{U}(x^k; G_k) - \tilde{U}(x^{k+1}; G_k) \ \ \textit{and} \ \ \tilde{k} = \arg\min_{k \in \{1, \ldots, N\}} \triangle \tilde{U}_k.$$

Then, we have

$$\mathsf{E}[\triangle U_{\tilde{k}}] \leq \mathsf{E}[\triangle \tilde{U}_{\tilde{k}}] \leq \left(\frac{2\,\sigma^q}{(\lambda\, p)^{q/p}} \sum_{k=1}^{N} \frac{1}{m_k^{q-1}} + \Phi(x^1) - \Phi^*\right) \Big/ N, \tag{24}$$

where $p \geq 2$ and $\frac{1}{p} + \frac{1}{q} = 1$. Moreover, if we assume that the batch sizes $m_k = m$ for $k = 1, \ldots, N$ with some $m \geq 1$, then

$$\mathsf{E}[\triangle U_{\tilde{k}}] \leq \mathsf{E}[\triangle \tilde{U}_{\tilde{k}}] \leq \frac{2\,\sigma^q}{(\lambda\,p)^{q/p}} \frac{1}{m^{q-1}} + \frac{\Phi(x^1) - \Phi^*}{N}. \tag{25}$$

Proof. Let $\delta_k = \nabla f(x^k) - G_k$ and denote

$$y^{k+1} = \arg\min_{y \in S} U(y; x^k) = \arg\min_{y \in S} \nabla f(x^k)^\top (y - x^k) + \frac{\lambda}{2}\|y - x^k\|_p^p + h(y). \tag{26}$$

Then it follows from Assumption 1.1 that

$$\begin{aligned}
\triangle \tilde{U}_k &= -G_k^\top (x^{k+1} - x^k) - \frac{\lambda}{2}\|x^{k+1} - x^k\|_p^p + h(x^k) - h(x^{k+1}) \\
&= \delta_k^\top (x^{k+1} - x^k) - \nabla f(x^k)^\top (x^{k+1} - x^k) - \frac{\lambda}{2}\|x^{k+1} - x^k\|_p^p + h(x^k) - h(x^{k+1}) \\
&\leq \delta_k^\top (y^{k+1} - x^k) + \delta_k^\top (x^{k+1} - y^{k+1}) + f(x^k) - f(x^{k+1}) + h(x^k) - h(x^{k+1}) \\
&\leq \delta_k^\top (y^{k+1} - x^k) + \|\delta_k\|_q \|x^{k+1} - y^{k+1}\|_p + \Phi(x^k) - \Phi(x^{k+1}) \\
&\leq \delta_k^\top (y^{k+1} - x^k) + \frac{2}{(\lambda\,p)^{q/p}} \|\delta_k\|_q^q + \Phi(x^k) - \Phi(x^{k+1}),
\end{aligned}$$

where the last inequality is due to Lemma 4.2 and the fact $\frac{1}{p} + \frac{1}{q} = 1$. Now, summing up the above inequalities for $k = 1, \ldots, N$, we obtain

$$\sum_{k=1}^N \triangle \tilde{U}_k \leq \sum_{k=1}^N \delta_k^\top (y^{k+1} - x^k) + \frac{2}{(\lambda\,p)^{q/p}} \sum_{k=1}^N \|\delta_k\|_q^q + \Phi(x^1) - \Phi^*. \tag{27}$$

Let $\xi_{[k-1]}$ be the random samples generated before iteration k. So, at the time x^k and y^{k+1} were determined, $\xi_{[k-1]}$ was already realized. Consequently, by Assumption (19), we have

$$\mathsf{E}\left[\delta_k^\top (y^{k+1} - x^k) \,\middle|\, \xi_{[k-1]}\right] = 0. \tag{28}$$

In addition, according to Assumption (20), one has

$$\mathsf{E}[\|\delta_k\|_q^q] = \frac{1}{m_k^q} \sum_{i=1}^{m_k} \mathsf{E}[\|\nabla f(x^k) - G(x^k, \xi_{k,i})\|_q^q] \leq \frac{1}{m_k^{q-1}} \sigma^q.$$

The above two formulas and (27) lead to

$$N\,\mathsf{E}[\triangle \tilde{U}_{\tilde{k}}] \leq \sum_{k=1}^N \mathsf{E}[\triangle \tilde{U}_k] \leq \frac{2\,\sigma^q}{(\lambda\,p)^{q/p}} \sum_{k=1}^N \frac{1}{m_k^{q-1}} + \Phi(x^1) - \Phi^*. \tag{29}$$

Furthermore, by definition of y^{k+1} in (26), one has

$$\begin{aligned}
\triangle U_k &= -\nabla f(x^k)^\top (y^{k+1} - x^k) - \frac{\lambda}{2}\|y^{k+1} - x^k\|_p^p + h(x^k) - h(y^{k+1}) \\
&= -\delta_k^\top (y^{k+1} - x^k) - G_k^\top (y^{k+1} - x^k) - \frac{\lambda}{2}\|y^{k+1} - x^k\|_p^p + h(x^k) - h(y^{k+1}) \\
&\leq -\delta_k^\top (y^{k+1} - x^k) + \triangle \tilde{U}_k,
\end{aligned}$$

which in combination with (28) and (29) yields (24). □

In order to apply Part (ii) of Lemma 2.2, we want $\mathsf{E}[\triangle U_{\tilde{k}}]$ to be upper bounded by $\frac{1}{2}\left(\frac{\epsilon}{\mathrm{diam}_p(S)\lambda^{1/p}}\right)^q$. Since the bound in the above theorem depends on m, we present the following corollary to show how m can be chosen so as to achieve the sharpest result.

Corollary 4.4 *Suppose $\triangle U_{\tilde{k}}$ is defined in Theorem 4.3, integer $p \geq 2$ and $\frac{1}{p}+\frac{1}{q}=1$. For a given sufficiently small ϵ, if the total number of calls $\bar{N}$ to the $\mathcal{SFO}$ is given by*

$$\bar{N} = \left\lceil \frac{(4\sigma)^p \lambda^{q-1} \mathrm{diam}_p(S)^{pq} (\Phi(x^1)-\Phi^*)}{p\,\epsilon^{pq}} \left(1+(q-1)^{-1/p+1/q}\right)^p \right\rceil, \tag{30}$$

and in each iteration of **Algorithm 3** *the batch size is set to be*

$$m = \left\lceil \min\left\{ \max\left\{ 1, \frac{\sigma((q-1)\bar{N})^{1/q}}{(\lambda p)^{1/p}(\Phi(x^1)-\Phi^*)^{1/q}} \right\}, \bar{N} \right\} \right\rceil,$$

then we have

$$\mathsf{E}[\triangle U_{\tilde{k}}] \leq \frac{1}{2}\left(\frac{\epsilon}{\mathrm{diam}_p(S)\lambda^{1/p}}\right)^q.$$

Proof. We first assume that

$$\bar{N} \geq \frac{\sigma^p (q-1)^{p/q}}{\lambda p (\Phi(x^1)-\Phi^*)^{p/q}} \tag{31}$$

or equivalently, $\bar{N} \geq \frac{\sigma((q-1)\bar{N})^{1/q}}{(\lambda p)^{1/p}(\Phi(x^1)-\Phi^*)^{1/q}}$. Since the batch size m_k at each iteration is identical, **Algorithm 3** can perform at most $N = \lfloor \bar{N}/m \rfloor$ iterations, which implies $N \geq \bar{N}/(2m)$. This fact together with (25) yields

$$\begin{aligned}
\mathsf{E}[\triangle U_{\tilde{k}}] &\leq \frac{2\,\sigma^q}{(\lambda p)^{q/p}} \frac{1}{m^{q-1}} + \frac{2m(\Phi(x^1)-\Phi^*)}{\bar{N}} \\
&\leq \frac{2\sigma(\Phi(x^1)-\Phi^*)^{1/p}}{(\lambda p\,(q-1)\,\bar{N})^{1/p}} + \frac{2(\Phi(x^1)-\Phi^*)}{\bar{N}} \left(1+\frac{\sigma((q-1)\bar{N})^{1/q}}{(\lambda p)^{1/p}(\Phi(x^1)-\Phi^*)^{1/q}}\right) \\
&= \frac{2(\Phi(x^1)-\Phi^*)}{\bar{N}} + \frac{2\sigma(\Phi(x^1)-\Phi^*)^{1/p}}{(\lambda p\,\bar{N})^{1/p}} \left((q-1)^{-1/p+1/q}\right),
\end{aligned}$$

where we used the fact $\frac{1}{p}+\frac{1}{q}=1$. If we further assume

$$\bar{N} \geq \frac{(\lambda p)^{q/p}\sigma^q}{\Phi(x^1)-\Phi^*}, \tag{32}$$

then

$$\mathsf{E}[\triangle U_{\tilde{k}}] \leq \frac{2\sigma(\Phi(x^1)-\Phi^*)^{1/p}}{(\lambda p\,\bar{N})^{1/p}} \left(1+(q-1)^{-1/p+1/q}\right).$$

Thus, by choosing $\bar{N}$ according to (30), when ϵ is sufficiently small so that (31) and (32) are satisfied, then the desired result follows.

□

4.2 Complexity for Smoothing Approximation

In this subsection, we shall consider a modified model where the nonsmooth part $h(x)$ is assumed to be concave rather than convex, while the smooth part $f(x)$ is still assumed to satisfy Assumption 1.1. In fact, this model is frequently encountered in several applications. For instance, in the L_2-L_q minimization problem [11], Assumption 1.1 holds with $p = 2$ since $f(x) = \|Ax - b\|_2^2$, and $h(x) = \|x\|_q^q$ with $0 < q < 1$ is a concave function. In this case, we apply the smoothing approximation to the nonsmooth function $h(x)$. In particular, consider the following convolution between h and μ:

$$h_\mu(x) := \int_{\mathbf{R}^n} h(x+y)\mu(y)dy = \mathsf{E}_\mu[h(x+Z)],$$

where Z is a random variable with probability density μ. It is well known (cf. [5]) that if μ is a density with respect to Lebesgue measure, then h_μ is differentiable. Furthermore, it is well known that (see, e.g., [14]) $h_\mu(x)$ and $h(x)$ can be bounded from each other using the properties of μ, as stated in the following lemma.

Lemma 4.5 *Let ξ be a random variable with the uniform density μ over the L_2-ball (radius r). Assume that $\sup\{\|g\|_2^2 \mid g \in \partial h(x)\} \le M$ for $x \in S + \mathcal{B}_2(0,r)$, where $\mathcal{B}_2(0,r) = \{y \mid \|y\|_2 \le r\}$ is the Euclidean ball with radius r. Let*

$$h_r(x) := \mathsf{E}_\mu[h(x + r\xi)]. \tag{33}$$

It holds that
(i) If $h(x)$ is concave then $h_r(x)$ is a concave function as well;
(ii) $h(x) \le h_r(x) \le h(x) + Mr$;
(iii) If $h(x)$ is differentiable, then $\mathsf{E}[\nabla h(x + r\xi)] = \nabla h_r(x)$ and $\mathsf{E}[\|\nabla h(x + r\,\xi) - \nabla h_r(x)\|_2^2] \le M^2$.

Now we consider the problem

$$\begin{array}{ll} \min & \Phi_r(x) := f(x) + h_r(x) \\ \text{s.t.} & x \in S \subseteq \mathbf{R}^n, \end{array} \tag{34}$$

Denote x^* and $\tilde{x}$ to be optimal solutions of (2) and (34) respectively. Then from Lemma 4.5 and the optimality of $\tilde{x}$, we have

$$f(\tilde{x}) + h(\tilde{x}) \le f(\tilde{x}) + h_r(\tilde{x}) \le f(x^*) + h_r(x^*) \le f(x^*) + h(x^*) + Mr, \tag{35}$$

which means that if the perturbation is small, then the smoothing version (34) is indeed a good approximation for the original problem (2). Denote

$$L_r(y; x^k) = (\nabla f(x^k) + \nabla h_r(x^k))^\top (x - x^k), \text{ and } \triangle L_k = L_r(x^k; x^k) - L_r(z^k; x^k),$$

where $z^k = \min_{z \in S} L_r(z; x^k)$. According to Lemma 2.2, x^k is an ϵ-stationary point of (34), if $\tilde{L}_k \le \epsilon$. We now propose a sampling-smoothing algorithm for (34) as follows:

Algorithm 4

Let $x^0 \in S$ be given and set $y^0 = x^0$.

for $k = 1, 2, \cdots, N$, **do**

Draw i.i.d. random samples $(\xi^{k,1}, \ldots, \xi^{k,m_k})$, and set $G_k = \frac{1}{m_k}\sum_{i=1}^{m_k} \nabla h(x^k + r\xi^{k,i})$.

Compute $y^k = \arg\min_{y\in S} \tilde{L}(y; x^k)$, where $\tilde{L}(y; G_k) = (\nabla f(x^k) + G_k)^\top (x - x^k)$.

Let $d^k = y^k - x^k$, and $\alpha_k = \arg\min_{\alpha\in[0,1]} \alpha\,(\nabla f(x^k) + G_k))^\top d^k + \alpha^p \frac{\lambda}{2}\|d^k\|_p^p$.

Set $x^{k+1} = (1 - \alpha_k)x^k + \alpha_k y^k$.

end for

The iteration complexity of this algorithm is presented in the following theorem.

Theorem 4.6 *Suppose* $\{x^k\}$ *is the sequence generated by* **Algorithm 4**. *Denote*

$$\triangle\tilde{L}_k = \tilde{L}(x^k; G_k) - \tilde{L}(y^k; G_k) \ \textit{and}\ \tilde{k} = \arg\min_{k\in\{1,\ldots,N\}} \triangle\tilde{L}_k.$$

For $\epsilon \le \mathrm{diam}_p^p(S)\lambda$, *let* $m_k = \left\lceil \frac{\mathrm{diam}_2^2(S)M^2N^2}{(\Phi(x^1)-\Phi^*)^2} \right\rceil$ *for* $k = 1, \ldots, N$ *and* $N = \left\lceil \frac{4(\Phi(x^1)-\Phi^*)(\mathrm{diam}_p^p(S)\lambda)^{q-1}}{\epsilon^q} \right\rceil$; *then we have*

$$\mathsf{E}[\triangle L_{\tilde{k}}] \le \mathsf{E}[\triangle\tilde{L}_{\tilde{k}}] \le \epsilon, \tag{36}$$

where $\frac{1}{p} + \frac{1}{q} = 1$.

Proof. Denote $\delta_k = \nabla h_r(x^k) - G_k$. By Assumption 1.1 and the optimality of α_k, and noting that $\frac{\epsilon}{\mathrm{diam}_p^p(S)\lambda} \le 1$ and $x^{k+1} - x^k = \alpha_k(y^k - x^k)$, we obtain the following sequence of inequalities

$$\begin{aligned}
& \left(\frac{\epsilon}{\mathrm{diam}_p^p(S)\lambda}\right)^{\frac{1}{p-1}} \triangle\tilde{L}_k - \delta_k^\top(x^{k+1} - x^k) - \frac{\lambda}{2}\left(\frac{\epsilon}{\lambda\mathrm{diam}_p(S)}\right)^{\frac{p}{p-1}} \\
\le\ & -\left(\frac{\epsilon}{\mathrm{diam}_p^p(S)\lambda}\right)^{\frac{1}{p-1}} (\nabla f(x^k) + G_k)^\top(y^k - x^k) - \frac{\lambda}{2}\frac{\|y^k - x^k\|_p^p}{\mathrm{diam}_p^p(S)}\left(\frac{\epsilon}{\lambda\mathrm{diam}_p(S)}\right)^{\frac{p}{p-1}} - \delta_k^\top(x^{k+1} - x^k) \\
\le\ & -\alpha_k(\nabla f(x^k) + G_k)^\top(y^k - x^k) - \frac{\lambda\alpha_k^p}{2}\|y^k - x^k\|_p^p - \delta_k^\top(x^{k+1} - x^k) \\
=\ & -(\nabla f(x^k) + G_k)^\top(\alpha_k(y^k - x^k)) - \frac{\lambda}{2}\|\alpha_k(y^k - x^k)\|_p^p - \delta_k^\top(x^{k+1} - x^k) \\
=\ & -(\nabla f(x^k) + G_k)^\top(x^{k+1} - x^k) - \frac{\lambda}{2}\|x^{k+1} - x^k\|_p^p - \delta_k^\top(x^{k+1} - x^k) \\
=\ & -(\nabla f(x^k) + \nabla h_r(x^k))^\top(x^{k+1} - x^k) - \frac{\lambda}{2}\|x^{k+1} - x^k\|_p^p \\
\le\ & f(x^k) - f(x^{k+1}) + h_r(x^k) - h_r(x^{k+1}) \\
=\ & \Phi_r(x^k) - \Phi_r(x^{k+1}),
\end{aligned}$$

where the last inequality follows from Assumption 1.1 and concavity of $h_r(\cdot)$. Dividing both sides of the above inequality by $\left(\frac{\epsilon}{\mathrm{diam}_p^p(S)\lambda}\right)^{\frac{1}{p-1}}$ and rearranging the terms yield

$$\triangle \tilde{L}_k \leq \left(\frac{\epsilon}{\mathrm{diam}_p^p(S)\lambda}\right)^{-\frac{1}{p-1}} \left(\Phi(x^k) - \Phi(x^{k+1}) + \|\delta_k\|_2 \|x^{k+1} - x^k\|_2\right) + \frac{\epsilon}{2}.$$

Since $h(\cdot)$ is concave and S is compact, $\bigcup_{x\in S+\mathcal{B}_2(0,r)} \partial h(x)$ is bounded. According to Lemma 4.5,

$$(\mathsf{E}[\|\delta_k\|_2])^2 \leq \mathsf{E}[\|\delta_k\|_2^2] = \frac{1}{m_k^2} \sum_{i=1}^{m_k} \mathsf{E}[\|\nabla h(x^k + r\xi^{k,i}) - \nabla h_r(x)\|_2^2] \leq \frac{M^2}{m_k}.$$

Therefore, summing over $k = 1, \ldots, N$ and taking expectation, one has

$$\begin{aligned}
N\mathsf{E}[\triangle \tilde{L}_{\tilde{k}}] &\leq \sum_{k=1}^{N} \mathsf{E}[\triangle \tilde{L}_k] \\
&\leq \left(\frac{\epsilon}{\mathrm{diam}_p^p(S)\lambda}\right)^{-\frac{1}{p-1}} \left(\Phi(x^1) - \Phi(x^{N+1}) + \sum_{k=1}^{N} \frac{M}{\sqrt{m_k}} \|x^{k+1} - x^k\|_2\right) + \frac{\epsilon}{2} N \\
&\leq \left(\frac{\epsilon}{\mathrm{diam}_p^p(S)\lambda}\right)^{-\frac{1}{p-1}} \left(\Phi(x^1) - \Phi^* + \mathrm{diam}_2(S) \sum_{k=1}^{N} \frac{M}{\sqrt{m_k}}\right) + \frac{\epsilon}{2} N.
\end{aligned}$$

When m_k and N are chosen as described, we have

$$\begin{aligned}
\mathsf{E}[\triangle \tilde{L}_{\tilde{k}}] &\leq \left(\frac{\epsilon}{\mathrm{diam}_p^p(S)\lambda}\right)^{-\frac{1}{p-1}} \left(\frac{\Phi(x^1) - \Phi^*}{N} + \frac{\mathrm{diam}_2(S) M}{\sqrt{m_1}}\right) + \frac{\epsilon}{2} \\
&\leq \left(\frac{\epsilon}{\mathrm{diam}_p^p(S)\lambda}\right)^{-\frac{1}{p-1}} \left(\frac{\Phi(x^1) - \Phi^*}{N} + \frac{\Phi(x^1) - \Phi^*}{N}\right) + \frac{\epsilon}{2} \\
&\leq \frac{\epsilon}{4} + \frac{\epsilon}{4} + \frac{\epsilon}{2} = \epsilon.
\end{aligned}$$

Since $z^k = \min_{z\in S} L_r(z; x^k)$, then

$$\begin{aligned}
\triangle L_{\tilde{k}} &= -(\nabla f(x^k) + \nabla h_r(x^k))^\top (z^k - x^k) \\
&= -(\nabla f(x^k) + G_k)^\top (z^k - x^k) - \delta_k^\top (z^k - x^k) \\
&\leq -(\nabla f(x^k) + G_k)^\top (y^k - x^k) - \delta_k^\top (z^k - x^k) = \triangle \tilde{L}_{\tilde{k}} - \delta_k^\top (z^k - x^k).
\end{aligned}$$

Let $\xi_{[k-1]}$ be the random samples generated before iteration k. The iterates x^k and z^k were determined after $\xi_{[k-1]}$ was realized. This fact combined with Lemma 4.5 implies that

$$\mathsf{E}\left[\delta_k^\top (z^k - x^k) \,\middle|\, \xi_{[k-1]}\right] = 0.$$

Therefore, $\mathsf{E}[\triangle L_{\tilde{k}}] \leq \mathsf{E}[\triangle \tilde{L}_{\tilde{k}}]$, and the theorem is proven. □

5 Iteration Complexity for Nonconvex Multi-block Optimization

In this section we consider the multi-block extension of (2):

$$\begin{array}{ll} \min & \Phi(x) := f(x_1, \cdots, x_d) + \sum_{i=1}^d h_i(x_i) \\ \text{s.t.} & x_i \in S_i \subseteq \mathbf{R}^{n_i}, \quad i = 1, \ldots, d, \end{array} \tag{37}$$

where f is differentiable but possibly nonconvex, and h_i is convex but possibly nonsmooth, $i = 1, \ldots, d$, and the feasible region S_i is convex and compact for all i (thus $S = \Pi_{i=1}^d S_i$ is compact as well). We denote

$$\mathrm{diam}_p(\overline{S}) = \max_{i\in\{1,\ldots,d\}} \max_{x_i, y_i \in S_i} \|x_i - y_i\|, \quad \text{and} \quad \mathrm{diam}_p(\underline{S}) = \min_{i\in\{1,\ldots,d\}} \max_{x_i, y_i \in S_i} \|x_i - y_i\|.$$

A well known technique for solving (37) is the so-called block coordinate descent (BCD) method. That is, at each iteration, a single block variable is optimized while all other blocks are fixed. In particular, at iteration k we solve the one-block problem exactly and denote

$$y_i^k \in \arg\min_{x_i \in S_i} f(x_1^{k-1}, \cdots, x_{i-1}^{k-1}, x_i, x_{i+1}^{k-1}, \cdots, x_d^{k-1}) + h_i(x_i), \quad i = 1, \ldots, d.$$

In the classical BCD method with Jacobian updating rule, the blocks are updated cyclicly by setting

$$x_i^k = y_i^k, \quad i = 1, \ldots, d.$$

Chen et al. [6] proposed another updating rule termed MBI (Maximum Block Improvement), where only the block with maximum improvement is updated at each step. Specifically, we first calculate the maximum improved block

$$i_0 \in \arg\max_{i\in\{1,\ldots,d\}} f(x^k) + h_i(x_i^k) - f(x_1^{k-1}, \cdots, x_{i-1}^{k-1}, y_i^k, x_{i+1}^{k-1}, \cdots, x_d^{k-1}) - h_i(y_i^k),$$

and then update the blocks by letting

$$x_{i_0}^k = y_{i_0}^k, \quad \text{and} \quad x_i^k = x_i^{k-1} \quad \text{for} \quad i \neq i_0.$$

To differentiate from the Jacobian style updating rule, a cyclic coordinate search is often referred as the BCD method of the Gauss-Seidel type, whose convergence under various settings has been established in [37, 39, 34]. Recently, the iteration complexity bounds were successfully established in some convex optimization problems [4, 22]. However, computational complexity analysis for *nonconvex* multi-block optimization is still very challenging. To the best of our knowledge, Dang and Lan [12] was probably the first paper to address this issue through a stochastic approximation method based on the approximation measure introduced in (5). Here we propose another method, based on the new notion that x is an ϵ-stationary point of (37) if

$$\nabla_i f(x)^\top (y_i - x_i) + h_i(y_i) - h(x_i) \geq -\epsilon \quad \forall y \in S_i, \quad i = 1, \ldots, d, \tag{38}$$

where $x = (x_1^\top, \cdots, x_d^\top)^\top$ and $y = (y_1^\top, \cdots, y_d^\top)^\top$.

In the following, we still assume that inequality (8) holds. Like in the single block-variables case as we discussed before, instead of solving the subproblems exactly, we shall use the following partially linearized (and p-powered upper bound) functions:

$$L_i(x_i; z) := f(z) + \nabla_i f(z)^\top (x_i - z_i) + h_i(x_i), \quad i = 1, \ldots, d, \tag{39}$$

$$U_i(x_i; z) := f(z) + \nabla_i f(z)^\top (x_i - z_i) + \frac{\lambda}{2}\|x_i - z_i\|_p^p + h_i(x_i), \quad i = 1, \ldots, d, \tag{40}$$

where $z = (z_1^\top, \cdots, z_d^\top)^\top$. Similar to Lemma 2.2, we have the following criteria to determine the ϵ-stationary point.

Lemma 5.1 *Given $\epsilon \geq 0$, for any $z \in S$,*
(i) if $\triangle_i L_z \leq \epsilon$ for $i = 1, \cdots, d$, then z is an ϵ-stationary point of (37);
(ii) if $\triangle_i U_z \leq \frac{1}{2}\left(\frac{\epsilon}{\mathrm{diam}_p(S_i)\lambda^{1/p}}\right)^q$ with $\frac{1}{p} + \frac{1}{q} = 1$ and $\epsilon \leq \mathrm{diam}_p^p(S_i)\lambda$, for $i = 1, \cdots, d$, then z is an ϵ-stationary point of (37).

Now we are ready to present our first algorithm for block optimization (37), where either the classical Jacobian updating rule or the MBI updating rule can be applied.

Algorithm 5

Let $x^0 \in S$ be given and set $y^0 = x^0$.
for $k = 1, 2, \cdots, N$, **do**
  for $i = 1, \cdots, d$, **do**
    $y_i^k = \arg\min_{y_i \in S_i} L_i(y_i; x^k)$, and let $d_i^k = y_i^k - x_i^k$;
    $\alpha_{k,i} = \arg\min_{\alpha \in [0,1]} \alpha \nabla_i f(x^k)^\top d_i^k + \alpha^p \frac{\lambda}{2}\|d_i^k\|_p^p + (1-\alpha)h_i(x_i^k) + \alpha\, h_i(y_i^k)$.
    Set $x_i^{k+1} = (1 - \alpha_{k,i})x_i^k + \alpha_{k,i} y_i^k$, when *Jacobian updating rule is applied.*
  end for
  (Or, calculate $i_0 \in \arg\max_{i \in \{1,\ldots,d\}} \triangle_i L_k$ and update $x_{i_0}^{k+1} = (1 - \alpha_{k,i_0})x_{i_0}^k + \alpha_{k,i_0} y_{i_0}^k$, $x_i^k = x_i^{k-1}$ if $i \neq i_0$ when the MBI updating rule is applied.)
end for

The computational complexity bound for this algorithm is established as follows:

Theorem 5.2 *For any $0 < \epsilon < \mathrm{diam}_p^p(\underline{S})\lambda$,* **Algorithm 5** *finds an ϵ-stationary point of (37) within $\left\lceil \frac{2(\mathrm{diam}_p(\overline{S})^p\lambda)^{q-1}(\Phi(x^1) - \Phi^*)}{\epsilon^q} \right\rceil$ steps.*

Proof. Since inequality (8) holds, one has

$$\begin{aligned}
&\sum_{i=1}^{d}\left(-\nabla_i f(x^k)^\top(x_i^{k+1}-x_i^k)-\frac{\lambda}{2}\|x_i^{k+1}-x_i^k\|_p^p+h_i(x_i^k)-h_i(x_i^{k+1})\right)\\
\leq\ & f(x^k)-f(x^{k+1})+\sum_{i=1}^{d}\left(h_i(x_i^k)-h_i(x_i^{k+1})\right)\\
=\ & \Phi(x^k)-\Phi(x^{k+1}). \qquad (41)
\end{aligned}$$

When the Jacobian updating rule is applied, by the definition of $\alpha_{k,i}$ we have

$$\alpha_{k,i}\left(\nabla_i f(x^k)^\top(y_i^k-x_i^k)+h_i(y_i^k)-h_i(x_i^k)\right)+\frac{\lambda\alpha_{k,i}^p}{2}\|y_i^k-x_i^k\|_p^p\leq 0.$$

Recall that $i_0\in\arg\max_{i\in\{1,\ldots,d\}}\triangle_i L_k$. Therefore,

$$\begin{aligned}
&-\alpha_{k,i_0}\left(\nabla_{i_0} f(x^k)^\top(y_{i_0}^k-x_{i_0}^k)+h_{i_0}(y_{i_0}^k)-h_{i_0}(x_{i_0}^k)\right)-\tfrac{\lambda\alpha_{k,i_0}^p}{2}\|y_{i_0}^k-x_{i_0}^k\|_p^p\\
\leq\ & \textstyle\sum_{i=1}^{d}\left(-\alpha_{k,i}\left(\nabla_i f(x^k)^\top(y_i^k-x_i^k)+h_i(y_i^k)-h_i(x_i^k)\right)-\tfrac{\lambda\alpha_{k,i}^p}{2}\|y_i^k-x_i^k\|_p^p\right)\\
=\ & \textstyle\sum_{i=1}^{d}\left(-\nabla_i f(x^k)^\top(\alpha_{k,i}(y_i^k-x_i^k))+h_i(x_i^k)-(1-\alpha_{k,i})h_i(x_i^k)-\alpha_{k,i}h_i(y_i^k)-\tfrac{\lambda}{2}\|\alpha_{k,i}(y_i^k-x_i^k)\|_p^p\right)\\
\leq\ & \textstyle\sum_{i=1}^{d}\left(-\nabla_i f(x^k)^\top(x_i^{k+1}-x_i^k)-\tfrac{\lambda}{2}\|x_i^{k+1}-x_i^k\|_p^p+h_i(x_i^k)-h_i(x_i^{k+1})\right),
\end{aligned} \qquad (42)$$

where the last inequality follows from the convexity of function $h_i(\cdot)$.

On the other hand, in the case of the MBI updating rule, we have

$$\begin{aligned}
&\sum_{i=1}^{d}\left(-\nabla_i f(x^k)^\top(x_i^{k+1}-x_i^k)-\frac{\lambda}{2}\|x_i^{k+1}-x_i^k\|_p^p+h_i(x_i^k)-h_i(x_i^{k+1})\right)\\
=\ & -\nabla_{i_0} f(x^k)^\top(x_{i_0}^{k+1}-x_{i_0}^k)-\frac{\lambda}{2}\|x_{i_0}^{k+1}-x_{i_0}^k\|_p^p+h_{i_0}(x_{i_0}^k)-h_{i_0}(x_{i_0}^{k+1}).
\end{aligned}$$

Thus, according to the convexity of $h_{i_0}(\cdot)$, inequality (42) still holds.

Therefore, in either cases, by (41), (42) and the optimality of α_{k,i_0} we have

$$\begin{aligned}
\left(\frac{\epsilon}{\mathrm{diam}_p^p(\overline{S})\lambda}\right)^{\frac{1}{p-1}}\max_i\triangle_i L_k &\leq \Phi(x^k)-\Phi(x^{k+1})+\frac{\epsilon^{\frac{p}{p-1}}\lambda\mathrm{diam}_p^p(S_{i_0})}{2(\mathrm{diam}_p^p(\overline{S})\lambda)^{\frac{p}{p-1}}}\\
&\leq \Phi(x^k)-\Phi(x^{k+1})+\frac{\epsilon^{\frac{p}{p-1}}}{2(\mathrm{diam}_p^p(\overline{S})\lambda)^{\frac{1}{p-1}}}.
\end{aligned}$$

Summing up the above inequality for $k = 1, \ldots, N$ yields

$$\begin{aligned}\left(\frac{\epsilon}{\mathrm{diam}_p^p(\overline{S})\lambda}\right)^{\frac{1}{p-1}} N \min_{k\in\{1,\ldots,N\}} \max_i \triangle_i L_k &\leq \left(\frac{\epsilon}{\mathrm{diam}_p^p(\overline{S})\lambda}\right)^{\frac{1}{p-1}} \sum_{k=1}^{N} \max_i \triangle_i L_k \\ &\leq \Phi(x^1) - \Phi(x^{N+1}) + N\frac{\epsilon^{\frac{p}{p-1}}}{2(\mathrm{diam}_p^p(\overline{S})\lambda)^{\frac{1}{p-1}}} \\ &\leq \Phi(x^1) - \Phi^* + N\frac{\epsilon^{\frac{p}{p-1}}}{2(\mathrm{diam}_p^p(\overline{S})\lambda)^{\frac{1}{p-1}}}.\end{aligned}$$

Thus, if

$$N = \left\lceil \frac{2(\mathrm{diam}_p(\overline{S})^p\lambda)^{q-1}(\Phi(x^1) - \Phi^*)}{\epsilon^q} \right\rceil,$$

then dividing both sides by $N\left(\frac{\epsilon}{\mathrm{diam}_p^p(\overline{S})\lambda}\right)^{\frac{1}{p-1}}$, we conclude that there must exist some $\tilde{k} \leq N$ such that

$$\triangle_i L_{\tilde{k}} \leq \epsilon, \quad \text{for all} \quad i = 1, \ldots, d,$$

which combined with Lemma 5.1 implies that $x^{\tilde{k}}$ is an ϵ-stationary point of (37). □

Our second algorithm to solve problem (37) uses an upper bound for the objective function. Again, we can either apply the classical Jacobian updating rule or the MBI updating rule.

Algorithm 6

Let $x^1 \in S$ be given

for $k = 1, 2, \cdots, N$, **do**

$y_i^{k+1} = \arg\min_{y_i \in S_i} U_i(y_i; x^k)$ for $i = 1, \ldots, d$.

Set $x_i^{k+1} = y_i^{k+1}$, $i = 1, \ldots, d$, when *Jacobian updating rule is applied.*

(Or, calculate $i_0 \in \arg\max_{i\in\{1,\ldots,d\}} \triangle_i U_k$ and update $x_{i_0}^{k+1} = y_{i_0}^{k+1}, x_i^{k+1} = x_i^k$ if $i \neq i_0$ when the *MBI* updating rule is applied.)

end for

Similarly, the computational complexity result for that algorithm can be established as follows.

Theorem 5.3 *For any* $0 < \epsilon < \mathrm{diam}_p^p(\underline{S})\lambda$, **Algorithm 6** *finds an* ϵ*-stationary point of* (37) *within* $\left\lceil \frac{2(\mathrm{diam}_p(\overline{S})^p\lambda)^{q-1}(\Phi(x^1)-\Phi^*)}{\epsilon^q} \right\rceil$ *steps.*

Similar arguments as in the proof of Theorem 5.2 can be used to prove the above theorem; we leave the details to the interested reader.

6 Numerical Experiments

In this section, we provide numerical performance of our algorithms for solving two nonconvex problems: the problem of finding the leading sparse principle component of tensor and the penalized zero-variance linear discriminant analysis.

6.1 Computing the Leading Sparse Principle Component of Tensor

The problem of finding the principle component (PC) that explains the most variance of a tensor $\mathcal{A}$ (with degree d) can be formulated as:

$$\begin{array}{ll} \min & \|\mathcal{A} - \lambda x_1 \otimes x_2 \otimes \cdots \otimes x_d\|_2^2 \\ \text{s.t.} & \lambda \in \mathbf{R},\ \|x_i\|_2^2 = 1, i = 1, 2, \ldots, d, \end{array}$$

where '$\otimes$' is the tensor outer-product operation. This problem bears different names including the tensor best rank-one approximation [24] and the Z-eigenvalue problem [32, 27], and various solution methods have been proposed: [6, 23, 24, 33].

Like in the matrix case, sparsity is desirable in tensor decomposition under various environments [1]. Consider the following sparse tensor PCA problem:

$$\begin{array}{ll} \min & \|\mathcal{A} - \lambda x_1 \otimes x_2 \otimes \cdots \otimes x_d\|_2^2 + \rho \sum_{i=1}^{d} \|x_i\|_0 \\ \text{s.t.} & \lambda \in \mathbf{R},\ \|x_i\|_2^2 = 1, i = 1, 2, \ldots, d, \end{array} \tag{43}$$

which is equivalent to

$$\begin{array}{ll} \min & -\mathcal{A}(x_1, x_2, \cdots, x_d) + \rho \sum_{i=1}^{d} \|x_i\|_0 \\ \text{s.t.} & \|x_i\|_2 \le 1, i = 1, 2, \ldots, d. \end{array}$$

To apply the algorithms discussed in the previous section, we replace $\|\cdot\|_0$ by $\|\cdot\|_1$, and arrive at the following formulation

$$\begin{array}{ll} \min & -\mathcal{A}(x_1, x_2, \cdots, x_d) + \rho \sum_{i=1}^{d} \|x_i\|_1 \\ \text{s.t.} & \|x_i\|_2 \le 1, i = 1, 2, \ldots, d. \end{array} \tag{44}$$

Denote the matrix

$$\mathcal{A}(x^{-ij}) := \mathcal{A}(x_1, \cdots, x_{j-1}, \cdot, x_{j+1}, \cdots, x_{i-1}, \cdot, x_{i+1}, \cdots, x_d),$$

and let $\tau = \max_{\|x\|_2 \le 1} \|\mathcal{A}(x^{-ij})\|_2$. Then for any $x = (x_1^\top, \cdots, x_d^\top)^\top, y = (y_1^\top, \cdots, y_d^\top)^\top$, and index

i, we have

$$\begin{aligned}
& \|\mathcal{A}(x_1,\cdots,x_{i-1},\cdot,x_{i+1},\cdots,x_d)-\mathcal{A}(y_1,\cdots,y_{i-1},\cdot,y_{i+1},\cdots,y_d)\|_2 \\
= & \left\|\sum_{j\neq i}\mathcal{A}(y_1,\cdots,y_{j-1},x_j,\cdots,x_{i-1},\cdot,x_{i+1},\cdots,x_d)-\mathcal{A}(y_1,\cdots,y_j,x_{j+1},\cdots,x_{i-1},\cdot,x_{i+1},\cdots,x_d)\right\|_2 \\
\leq & \sum_{j\neq i}\|\mathcal{A}(y_1,\cdots,y_{j-1},x_j,\cdots,x_{i-1},\cdot,x_{i+1},\cdots,x_d)-\mathcal{A}(y_1,\cdots,y_j,x_{j+1},\cdots,x_{i-1},\cdot,x_{i+1},\cdots,x_d)\|_2 \\
\leq & \tau\sum_{j\neq i}\|x_j-y_j\|_2.
\end{aligned}$$

Consequently,

$$\|\nabla\mathcal{A}(x)-\nabla\mathcal{A}(y)\|_2\leq\tau(d-1)\sum_{i=1}^{d}\|x_i-y_i\|_2\leq\tau d(d-1)\|x-y\|_2,$$

which means that Assumption 1.1 holds for $p=q=2$. Therefore our **Algorithm 5** and **Algorithm 6** can be applied to solve problem (44). When **Algorithm 5** is applied, the subproblem is in the form of

$$\min_{\|y\|_2^2\leq 1}\{-y^\top b+\rho\|y\|_1\}. \tag{45}$$

Denote $z(j)=\mathrm{sign}(b(j))\max\{|b(j)|-\rho,0\}\ \forall j$. Problem (45) has a closed form solution

$$y^*=\begin{cases} z/\|z\|_2, & \text{if } \|z\|_2\neq 0\\ 0, & \text{otherwise.}\end{cases}$$

In **Algorithm 6** the subproblem under consideration is given by $\min_{\|y\|_2^2\leq 1}\{-y^\top b+\rho\|y\|_1+\frac{\lambda}{2}\|y\|_2^2\}$, which has a closed form solution $y^*=\frac{z}{\lambda+\max\{0,\|z\|_2-\lambda\}}$.

One undesirable property of the relaxed formulation (44) is that we may possibly get a zero solution; i.e. $x_i=0$ for some i, which leads to $\mathcal{A}(x_1,x_2,\cdots,x_d)=0$. To prevent this from happening, we also apply the BCD method with the Jacobian updating rule to the following equality constraint problem:

$$\begin{aligned}
\min\quad & -\mathcal{A}(x_1,x_2,\cdots,x_d)+\rho\sum_{i=1}^{d}\|x_i\|_1 \\
\text{s.t.}\quad & \|x_i\|_2=1, i=1,2,\ldots,d,
\end{aligned} \tag{46}$$

and compare the results with those returned by our proposed algorithms in Table 1.

In the tests, we let $\lambda=20$, $\rho=0.85$, and set the maximum iteration number to be 2000; we only apply the Jacobian updating rule in the implementation of **Algorithm 5** and **Algorithm 6**. For each fixed dimension, we randomly generate 10 instances which are the fourth order tensors and the corresponding problems are solved by the three methods, starting from the same initial point. In Table 1, 'Val.' refers to the value $\mathcal{A}(x_1,x_2,\cdots,x_d)$. From this table, we see that **Algorithm 5** is the most stable method for the sparse tensor PCA problem, as it is able to find a nonzero local

minimum within a few hundred steps in most cases, with reasonably sparse solutions. For the same collection of instances, **Algorithm 6** falls into the zero solution in many cases, while the BCD method for (46) on the other hand, ends up in a local minimum point instead of local maximum for quite a few instances.

6.2 Penalized Zero-Variance Linear Discriminant Analysis

The penalized version of the zero-variance discriminant analysis is presented by Ames and Hong [2], aiming to perform linear discriminant analysis and feature selection on high-dimensional data simultaneously. To be more specific, the problem under consideration can be described as follows:

$$\min_{x\in\mathbf{R}^n,\, x^\top x\leq 1} -\frac{1}{2}x^\top N^\top BNx + \gamma\sum_{i=1}^m \sigma_i|(DNX)_i|, \tag{47}$$

where $B \in \mathbf{R}^{m\times m}$ is a positive semidefinite matrix, $D \in \mathbf{R}^{m\times m}$ and $N \in \mathbf{R}^{m\times n}$ are orthogonal matrices. Thus, the objective is the summation of a concave function and a convex function. Due to Corollary 3.2, **Algorithm 1** with $\alpha_k = 1$ for all k can achieve $O(1/\epsilon)$ iteration complexity. In particular, at the k-th iteration, based on point x^k we can find the next point x^{k+1} by optimizing a homogeneous convex problem:

$$x^{k+1} := \arg_{x\in\mathbf{R}^n} -(x^k)^\top N^\top BNx + \gamma\sum_{i=1}^m \sigma_i|(DNX)_i|,$$

which can be solved by CVX [20] efficiently. While Ames and Hong [2] proposes *alternating direction method of multipliers* (ADMM) to solve (47). The advantage of that approach is that the subproblems have closed form solutions.

For fixed dimension, we compare our approach with ADMM in [2] for 10 randomly generated instances, and the results are provided in Table 2. It appears that these two methods produce sequences that converge to the same point, but **Algorithm 1** costs much less iterations.

References

[1] G.I. Allen, *Sparse Higher-Order Principal Components Analysis*, Proceedings of the 15th International Conference on Artificial Intelligence and Statistics, 2012. 2, 24

[2] B.P.W. Ames and M. Hong, *Alternating Direction Method of Multipliers for Sparse Zero-Variance Discriminat Analysis and Principal Component Analysis* arXiv:1401.5492, 2014. 26

[3] A. Beck and M. Teboulle, *A Fast Iterative Shrinkage-Thresholding Algorithm for Linear Inverse Problems*, SIAM Journal on Imaging Sciences, 2, 183-202, 2009. 5

[4] A. Beck and L. Tetruashvili, *On the convergence of block coordinate descent type methods*, SIAM Journal on Optimization, 23, 2037-2060, 2013. 20

[5] D.P. Bertsekas, Stochastic Optimization Problems with Nondifferentiable Cost Functional, Journal of Optimization Theory and Application, 12, 218-231, 1973 17

[6] B. Chen, S. He, Z. Li, and S. Zhang, *Maximum Block Improvement and Polynomial Optimization*, SIAM Journal on Optimization, 22, 87-107, 2012. 20, 24

[7] C. Cartis, N.I.M. Gould, and Ph.L. Toint, *On the complexity of steepest descent, Newton's and regularized Newton's methods for nonconvex unconstrained optimization*, SIAM Journal on Optimization, 20(6), 2833-2852, 2010. 2

[8] C. Cartis, N.I.M. Gould, and Ph.L. Toint, *Adaptive cubic overestimation methods for unconstrained optimization. Part II: worst-case function-evaluation complexity*, Mathematical Programming, Series A, 130(2), 295-319, 2011. 2

[9] C. Cartis, N.I.M. Gould, and Ph.L. Toint, *An adaptive cubic regularization algorithm for nonconvex optimization with convex constraints and its function-evaluation complexity*, IMA Journal of Numerical Analysis, (to appear), 2013. 2, 3, 5

[10] C. Cartis, N.I.M. Gould, and Ph.L. Toint, On the complexity of finding first-order critical points in constrained nonlinear optimization, Mathematical Programming Series A, (to appear), 2013. 2

[11] X. Chen, D. Ge, Z. Wang, and Y. Ye, *Complexity of Unconstrained $L_2 - L_p$ Minimization*, Mathematical Programming, 143, 371-383, 2014. 17

[12] C.D. Dang and G. Lan, *Stochastic Block Mirror Descent Methods for Nonsmooth and Stochastic Optimization*, arXiv:1309.2249, 2013 20

[13] O. Devolder, G. François and Yu. Nesterov, *First-Order Methods of Smooth Convex Optimization with Inexact Oracle*, Mathematical Programming Series A, (to appear), 2013. 4

[14] J.C. Duchi, P.L. Bartlett, and M.J. Wainwright, *Randomized Smoothing for Stochastic Optimization*, SIAM Journal on Optimization, 22, 674-701, 2012. 17

[15] M. Frank and P. Wolfe, *An Algorithm for Quadratic Programming, Naval Research Logistics Quarterly*, 3, 95-110, 1956. 5

[16] R.M. Freund and P. Grigas, *New Analysis and Results for the Conditional Gradient Method*, arXiv:1307.0873, 2013. 5

[17] D. Ge, R. He, and S. He, *A Three Criteria Algorithm for $L_2 - L_p$ Minimization Problem With Linear Constraints*, Techniqual Report, 2014. 4, 8

[18] S. Ghadimi, G. Lan, and H. Zhang, *Mini-batch Stochastic Approximation Methods for Nonconvex Stochastic Composite Optimization*, arXiv:1308.6594, 2013. 2, 3, 13

[19] S. Ghadimi and G. Lan, *Stochastic First- and Zeroth-order Methods for Nonconvex Stochastic Programming*, SIAM Journal on Optimization, 23(4), 2341-2368, 2013. 2, 13

[20] M. Grant and S. Boyd, *CVX: Matlab software for disciplined convex programming, version 1.21*, http://cvxr.com/cvx, May 2010. 26

[21] Z. Harchaoui, A. Juditsky, and A.S. Nemirovski, *Conditional Gradient Algorithm for Machine Learing*, NIPS OPT workshop, 2012.

[22] M. Hong, X. Wang, M. Razaviyayn, and Z.-Q. Luo, *Iteration complexity analysis of block coordinate descent methods*, arXiv:1310.6957, 2013. 20

[23] B. Jiang, S. Ma, and S. Zhang, *Tensor Principal Component Analysis via Convex Optimization*, Mathematical Programming Series A, (to appear), 2014. 24

[24] E. Kofidis and P.A. Regalia, *On the best rank-1 approximation of higher-order supersymmetric tensors*, SIAM Journal on Matrix Analysis and Applications, 23, 863C884, 2002. 24

[25] G. Lan, *An optimal method for stochastic composite optimization*, Mathematical Programming, 133, 365-397, 2012. 13

[26] G. Lan, *The Complexity of Large-scale Convex Programming under a Linear Optimization Oracle*, arXiv:1309.5550, 2013. 5

[27] L.-H. Lim, *Singular Values and Eigenvalues of Tensors: A Variational Approach*, Proceedings of the IEEE International Workshop on Computational Advances in Multi-Sensor Adaptive Processing, 1, 129-132, 2005. 24

[28] Y. Liu, S. Ma, Y. Dai and S. Zhang, *A Smoothing SQP Framework for a Class of Composite L_q Minimization over Polyhedron*, 2014. http://www.optimization-online.org/DB_HTML/2014/07/4467.html. 5, 6

[29] A.S. Nemirovski, A. Juditsky, G. Lan, and A. Shapiro, *Robust Stochastic Approximation Approach to Stochastic Programming*, SIAM Journal on Optimization, 19, 1574-1609, 2009. 13

[30] A.S. Nemirovski and D. Yudin, *Problem Complexity and Method Efficiency in Optimization*, Wiley-Interscience Series in Discrete Mathematics, John Wiley, XV, 1983. 13

[31] Yu. Nesterov, *Introductory Lectures on Convex Optimization. Applied Optimization*, Kluwer Academica Publishers, Dordrecht, The Netherlands, 2004. 2

[32] L. Qi, *Eigenvalues of a Real Supersymmetric Tensor*, Journal of Symbolic Computation, 40, 1302-1324, 2005. 24

[33] L. Qi, F. Wang, and Y. Wang, *Z-eigenvalue methods for a global polynomial optimization problem*, Mathematical Programming, Series A, 118, 301-316, 2009. 24

[34] M. Razaviyayn, M. Hong, and Z.-Q. Luo, *A Unified Convergence Analysis of Block Successive Minimization Methods for Nonsmooth Optimization*, SIAM Journal on Optimization, 23, 1126-1153, 2013. 20

[35] H. Robbins and S. Monro, *A stochastic approximation method*, Annals of Mathematical Statistics, 22, 400-407, 1951. 12

[36] A. Saha and A. Tewari, *On the Nonasymptotic Convergence of Cyclic Coordinate Descent Methods*, SIAM Journal on Optimization, 23, 576-601, 2013.

[37] P. Tseng, *Convergence of a Block Coordinate Descent Method for Nondifferentiable Minimization*, Journal of Optimization Thoery and Applications, 103, 475-494, 2001. 20

[38] X. Wang, S. Ma, and Y.-X. Yuan, *Penalty Methods with Stochastic Approximation for Stochastic Nonlinear Programming*, arXiv:1312.2690, 2013. 13

[39] Y. Xu and W. Yin, *A Block Coordinate Descent Method for Regularized Multiconvex Optimization with Applications to Nonnegative Tensor Factorization and Completion*, SIAM Journal on Imaging Sciences 6, 1758-1789, 2013 20

Inst. #	BCD			Algorithm 6			Algorithm 5		
	Val.	$\sum_{i=1}^{d} \|x_i\|_0$	Iter.	Val.	$\sum_{i=1}^{d} \|x_i\|_0$	Iter.	Val.	$\sum_{i=1}^{d} \|x_i\|_0$	Iter.
Dimension $n = 8$									
1	4.99	20	2000	7.76	21	99	6.76	18	92
2	6.42	14	96	0.00	0	20	0.00	8	13
3	-6.41	16	46	6.56	16	213	6.56	16	139
4	-7.16	17	83	0.00	0	24	0.00	16	15
5	-6.36	18	99	6.78	18	185	6.19	15	91
6	-8.36	20	132	8.98	22	156	8.98	22	148
7	6.10	18	2000	7.51	18	154	7.51	18	98
8	-7.14	20	95	7.14	20	357	7.14	20	295
9	8.73	21	127	8.50	21	227	8.50	21	155
10	6.25	16	327	7.41	19	207	7.41	19	119
Dimension $n = 12$									
1	-7.80	20	2000	8.34	27	841	8.23	23	166
2	-9.52	31	293	8.16	23	222	8.73	26	315
3	9.17	28	282	10.19	33	528	10.19	33	362
4	8.50	22	257	7.63	21	230	8.74	22	296
5	-9.58	22	2000	8.61	23	314	8.61	24	221
6	9.95	28	267	8.48	22	211	8.48	22	151
7	8.88	23	142	0.00	0	18	0.00	24	11
8	-8.42	27	263	8.55	27	250	8.55	27	154
9	-8.64	26	2000	8.81	24	166	8.49	31	53
10	9.54	30	208	8.89	26	131	8.89	26	97
Dimension $n = 20$									
1	6.80	39	2000	0.00	0	31	11.52	41	290
2	-12.95	49	278	0.00	0	17	11.77	44	112
3	-11.44	38	277	11.08	39	156	12.71	42	195
4	-11.22	40	766	0.00	0	20	11.50	39	141
5	11.51	38	1267	0.00	0	18	0.00	0	11
6	12.42	44	808	11.44	36	225	11.47	36	144
7	-12.22	47	2000	11.28	34	269	13.20	49	241
8	11.35	39	2000	11.03	35	191	10.80	40	211
9	11.74	44	2000	11.88	37	194	12.34	47	199
10	-11.49	46	493	11.60	42	172	11.79	43	454
Dimension $n = 30$									
1	14.13	52	1673	0.00	0	12	15.22	58	311
2	0.82	41	2000	0.00	0	25	14.36	53	214
3	-15.23	59	1589	0.00	0	15	14.28	55	238
4	0.25	59	2000	0.00	0	15	0.00	0	9
5	0.74	43	2000	14.36	56	175	14.36	56	101
6	0.53	35	2000	13.33	47	454	15.24	58	267
7	0.48	43	2000	13.19	51	377	13.39	51	162
8	0.38	40	2000	0.00	0	14	0.00	0	9
9	-13.32	48	1523	13.93	56	692	12.36	47	374
10	11.84	40	1034	13.75	51	455	15.18	61	361

Table 1: Numerical results for sparse tensor PCA problem.

Inst. #	Algorithm 1		ADMM	
	Obj.Val.	Iter.	Obj.Val.	Iter.
Dimension $n = 50$, $m = 100$				
1	-279.372	93	-279.370	437
2	-273.868	61	-273.868	739
3	-263.741	72	-263.743	426
4	-257.462	82	-257.465	566
5	-263.832	114	-263.848	652
6	-291.784	42	-291.784	220
7	-277.502	46	-277.503	250
8	-279.291	65	-279.297	343
9	-265.733	145	-265.741	1029
10	-269.525	62	-269.523	273
Dimension $n = 100$, $m = 200$				
1	-574.195	133	-574.199	728
2	-553.266	78	-553.270	587
3	-564.459	161	-564.456	648
4	-586.036	46	-586.039	240
5	-554.430	135	-554.447	880
6	-563.877	59	-563.883	308
7	-559.835	72	-559.844	434
8	-557.033	157	-557.042	823
9	-558.926	155	-558.929	674
10	-573.841	198	-573.861	1019
Dimension $n = 200$, $m = 400$				
1	-1172.918	85	-1172.916	638
2	-1154.821	168	-1154.834	993
3	-1139.545	129	-1139.554	662
4	-1121.581	245	-1121.587	922
5	-1176.465	116	-1176.465	640
6	-1149.466	106	-1149.472	521
7	-1151.241	90	-1151.245	417
8	-1156.047	339	-1156.068	1533
9	-1135.454	59	-1135.518	474
10	-1140.661	512	-1140.704	2840

Table 2: Numerical results for Penalized Zero-Variance Linear Discriminant Analysis.